arXiv:1605.07785v1 [cs.LG] 25 May 2016

# Geometry-aware Stationary Subspace Analysis

Inbal Horev[1], Florian Yger[2], and Masashi Sugiyama[1]

[1]Graduate School of Frontier Sciences, University of Tokyo
[2]LAMSADE Université Paris-Dauphine

**Abstract**

In many real-world applications observed data exhibits *non-stationarity*, i.e., its distribution changes over time. One approach to handling non-stationarity is to remove or minimize it before attempting to analyze the data. In the context of brain computer interface (BCI) data analysis this may be done by means of *stationary subspace analysis* (SSA). The SSA method finds a matrix that projects the data onto a stationary subspace by optimizing a cost function based on a matrix divergence. In this work we present an alternative method for SSA based on a symmetrized version of this matrix divergence. We show that doing so frames the problem in terms of distances between symmetric positive definite (SPD) matrices, suggesting a geometric interpretation of the problem. Stemming from this geometric viewpoint, we introduce and analyze a method which utilizes the geometry of the SPD matrix manifold and the invariance properties of its metrics. We demonstrate the usefulness of our method in experiments on both synthesized and real-world data.

## 1 Introduction

A common assumption in statistical modeling is that the distribution of observed data does not change over time, i.e., that it is *stationary*. In most cases it is this assumption of stationarity which allows results to be effectively generalized from the sample to the population. When the stationarity assumption is violated, as is often the case in real-world applications such as speech enhancement [11] or neurological data analysis [25], specialized machine learning methods must be developed in order to maintain adequate prediction capabilities.

A relatively well-studied non-stationary setting is covariate-shift [26], in which the input distribution changes but the conditional distribution of the outputs does not. The problem of covariate-shift has received growing attention in recent years, and many theoretical and practical aspects have been addressed (see [28] for an in-depth exploration of this topic).

These works typically do not aim to remove or reduce non-stationarity in the data, but rather they try to cope with its existence. A different approach

is to remove, or minimize, any existing non-stationarities before attempting to analyze the collected data. In the context of brain computer interface (BCI) data analysis, two such note-worthy methods are stationary subspace analysis (SSA) [30] and stationary common spacial patterns (sCSP) [32].

Similar in spirit to independent component analysis (ICA) [16], SSA statistically models the data as a mixture of stationary and non-stationary signals. Unlike ICA, however, the signals are not assumed to be independent of each other. The data is first split into (possibly overlapping) time frames called epochs. Then a projection matrix is found by optimizing a cost function based on the divergence between distributions in various epochs.

The other method, by now quite a standard step for classification tasks in BCI systems, is the (supervised) sCSP method. Its goal is to project the data onto a subspace in which the various data classes are more separable. The sCSP method directs this subspace towards a stationary subspace by means of regularization.

Although SSA is essentially an unsupervised method, variations of it exist which are useful for supervised tasks such as classification [25]. These methods attempt to remove non-stationarity while keeping the discriminative inter-class variations intact. In this work we present an alternative method for stationary subspace extraction, focusing for the moment on the unsupervised setting.

Unlike SSA, the inputs to our method are not the raw signals themselves, but rather covariance matrices, computed from the signals (or from features based on the signals) using one of the many existing covariance estimators (e.g., [20]). Covariance matrices have gained increasing attention in recent years, and are now commonly used in many machine learning and signal processing applications such as computer vision applications [29], brain imaging [24] and BCI data analysis [6]. Their rich mathematical structure has been extensively studied [7], and advances in optimization methods on matrix manifolds in recent years have motivated the development of geometric methods for various machine learning tasks such as dictionary learning [9], metric learning [19] and dimensionality reduction [13]. This in turn motivates a covariance-based approach which utilizes the geometric properties of the symmetric positive definite (SPD) matrix manifold for SSA.

## 2 Geometry-aware stationary subspace analysis

As discussed in the introduction, the task of extracting the stationary part from an observed mixture of stationary and non-stationary signals is essential in various applications. In this section we present our approach to this problem. We name it *geometry-aware SSA* (gaSSA) since we utilize the geometric properties of covariance matrices. To find a stationary subspace, SSA uses a cost function which is based on a matrix divergence. As a first step, we suggest a formulation that uses a symmetrized version of the same matrix divergence. We then show that this symmetrized matrix divergence is a distance between SPD matrices and offer a geometric interpretation to the problem of SSA.

We begin with a formal statement of the problem. To this end we provide

a review of the original SSA model and framework [30]. Next, we present the symmetrized matrix divergence and discuss its relation to the Riemannian geometry of SPD matrices. Finally, we end this section by introducing and analyzing a general formulation for our geometry-aware SSA.

## 2.1 Stationary subspace analysis

Let $x(t) \in \mathbb{R}^D$ be a vector of $D$ input signals, composed of $m$ stationary sources $s^{\mathfrak{s}}(t) = [s_1(t), \ldots, s_m(t)]^\top$ ($\mathfrak{s}$-sources) and $D-m$ non-stationary sources $s^{\mathfrak{n}}(t) = [s_{m+1}(t), \ldots, s_D(t)]^\top$ ($\mathfrak{n}$-sources), mixed by a linear mixing transformation,

$$x(t) = As(t) = \left[A^{\mathfrak{s}} A^{\mathfrak{n}}\right] \begin{bmatrix} s^{\mathfrak{s}}(t) \\ s^{\mathfrak{n}}(t) \end{bmatrix}, \tag{1}$$

where $A \in GL_D(\mathbb{R})$, the general linear group of size $D$ over $\mathbb{R}$, i.e., the set of all $D \times D$ invertible matrices with entries in $\mathbb{R}$. The spaces spanned by the column vectors $A^{\mathfrak{s}}$ and $A^{\mathfrak{n}}$ are referred to as the $\mathfrak{s}$-space and $\mathfrak{n}$-space, respectively.

The SSA model makes relatively few assumptions on the $\mathfrak{s}$- and $\mathfrak{n}$-sources. First, the $\mathfrak{s}$-sources are stationary only in the weak (or wide) sense [17]. That is, their first and second moments are required to be constant in time. For the $\mathfrak{n}$-sources, their first two moments may vary between epochs of length $T$ denoted by $\tau_i = [t_0(i), \ldots, t_0(i)+T]$, where $t_0(i)$ is the start time of the $i$-th epoch. The sources do not necessarily follow a Gaussian distribution, but non-stationarities are assumed to be visible in the first two moments. Furthermore, this model does not assume that the sources are independent, namely, their covariance matrix is given by

$$\Sigma(\tau_i) = \mathbb{E}\left[s(\tau_i)s(\tau_i)^\top\right] = \begin{pmatrix} \Sigma^{\mathfrak{s}} & \Sigma^{\mathfrak{sn}}(\tau_i) \\ \Sigma^{\mathfrak{sn}}(\tau_i)^\top & \Sigma^{\mathfrak{n}}(\tau_i) \end{pmatrix}, \tag{2}$$

where $\Sigma^{\mathfrak{s}} \in \mathbb{R}^{m \times m}$, $\Sigma^{\mathfrak{n}} \in \mathbb{R}^{(D-m)\times(D-m)}$ and $\Sigma^{\mathfrak{sn}} \in \mathbb{R}^{m\times(D-m)}$. Note that since $\Sigma^{\mathfrak{s}}$ is time independent we have dropped the notation $(\tau_i)$.

The goal of SSA is to find a de-mixing transformation $\widehat{A}^{-1}$ that separates the $\mathfrak{s}$-sources from the $\mathfrak{n}$-sources. This matrix $\widehat{A}^{-1}$ is not unique, but rather undetermined up to scaling, sign and linear transformations within the $\mathfrak{s}$- and $\mathfrak{n}$-spaces. So, without loss of generality, the data may be centered and whitened such that the $\mathfrak{s}$-sources have a zero mean and a diagonal covariance matrix with unit variance.[1] Put differently, the de-mixing matrix is written as $\widehat{A}^{-1} = \widehat{B}Z$ where $Z = \mathrm{Cov}(x)^{-1/2}$ is a whitening matrix created by a covariance estimator $\mathrm{Cov}(\cdot)$ (in this case it is the empirical estimator) and $\widehat{B} \in \mathcal{O}_D = \left\{V \in \mathbb{R}^{D\times D} \,:\, V^\top V = \mathbb{I}\right\}$, the set of all $D \times D$ orthogonal matrices.

To find the matrix $\widehat{B}$, the signals are split into $N$ epochs $\tau_1, \ldots, \tau_N$ of length $T$. Each epoch is characterized by its empirical mean $\widehat{\mu}_i$ and covariance $\widehat{\Sigma}_i$. Then for each epoch, the mean and covariance of the $\mathfrak{s}$-sources may be written as

$$\widehat{\mu}_i^{\mathfrak{s}} = \mathbb{I}_D^m \widehat{B} Z \widehat{\mu}_i, \qquad \widehat{\Sigma}_i^{\mathfrak{s}} = \mathbb{I}_D^m \widehat{B} Z \widehat{\Sigma}_i \left(\mathbb{I}_D^m \widehat{B} Z\right)^\top, \tag{3}$$

[1] This is also common practice in ICA [16]

where $\mathbb{I}_D^m$ is the $D \times D$ identity matrix, truncated to the first $m$ columns. Since the true $\mu^{\mathfrak{s}}$ and $\Sigma^{\mathfrak{s}}$ are by definition stationary, the matrix $\widehat{B}$ is that for which $\widehat{\mu}_i^{\mathfrak{s}}$ and $\widehat{\Sigma}_i^{\mathfrak{s}}$ vary the least across all epochs. Owing to the maximum entropy principle, SSA uses the Kullback-Leibler (KL) divergence between Gaussian distributions to compare the epoch distributions up to their second moment. The matrix $\widehat{B}$ is thus found by minimizing the following cost function:

$$\mathcal{L}\left(\widehat{B}\right) = \sum_{i=1}^{N} D_{\mathrm{KL}}\left[\mathcal{N}\left(\widehat{\mu}_i^{\mathfrak{s}}, \widehat{\Sigma}_i^{\mathfrak{s}}\right) \,||\, \mathcal{N}(0, \mathbb{I})\right] = -\sum_{i=1}^{N}\left(\log\det\widehat{\Sigma}_i^{\mathfrak{s}} + \widehat{\mu}_i^{\mathfrak{s}\top}\widehat{\mu}_i^{\mathfrak{s}}\right), \quad (4)$$

where $D_{\mathrm{KL}}$ is the KL divergence and $\mathcal{N}(\mu, \Sigma)$ denotes a multivariate Gaussian distribution with mean $\mu$ and covariance $\Sigma$.

## 2.2 Symmetrized matrix divergence

SSA and its variants use in their cost function the KL divergence between Gaussian distributions. In what follows we assume that these distributions have a zero mean.Under this assumption the KL divergence is a Bregman matrix divergence [4] between covariance matrices. The family of Bregman matrix divergences is generally defined as

$$D_\Phi(X, Y) = \Phi(X) - \Phi(Y) - \mathrm{tr}\left(\left(\nabla\Phi(Y)\right)^\top (X - Y)\right). \quad (5)$$

$D_{\mathrm{KL}}$ is obtained for $\Phi(X) = -\log\det X$, so it is often called the log-determinant divergence [18].

Bregman matrix divergences are useful in machine learning and have a number of useful properties [18], such as linearity and convexity in the first argument (and, in the case of the KL divergence, also in the second). However, as can be seen from their definition, they are asymmetric and do not satisfy the triangle equality. In particular, in our case we have that $D_{\mathrm{KL}}(X \,||\, Y) \neq D_{\mathrm{KL}}(Y \,||\, X)$ for two arbitrary matrices $X \neq Y$. Subsequently, symmetrized versions of the Bregman matrix divergence, in particular Jensen-Bregman divergences, have been studied in recent years [23]. In the case of the KL divergence (for zero mean distributions), this gives

$$\begin{aligned} D_{\mathrm{JBLD}}(X, Y) &= \frac{1}{2}\left[D_{\mathrm{KL}}\left(X \,||\, \frac{1}{2}(X+Y)\right) + D_{\mathrm{KL}}\left(Y \,||\, \frac{1}{2}(X+Y)\right)\right] \\ &= \log\det\left(\frac{1}{2}(X+Y)\right) - \frac{1}{2}\log\det(XY) \end{aligned} \quad (6)$$

and is called the Jensen-Bregman log-determinant (JBLD) divergence [10].

The JBLD has many favorable properties (see [27]), primarily that its square root comprises a metric on the SPD matrix manifold. Moreover, in Sra [27] it has been shown that it is a close approximation to the *affine invariant Riemannian metric* (AIRM) [7] and shares many of its mathematical properties. The practical properties of both metrics, in particular their invariance properties, will be

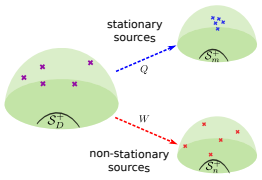


Figure 1: An illustration of our covariance-based approach. A set of input covariance matrices $\{\Sigma_i \in \mathcal{S}_D^+\}$ are made up of a mixture (purple) of stationary (blue) and non-stationary (red) parts. Due to non-stationarities, in the original space the matrices are spread out. The matrices are mapped to two lower dimensional spaces, namely the $\mathfrak{s}-$ (stationary) and $\mathfrak{n}-$ (non-stationary) spaces, via the matrices $Q$ and $W$, respectively. In the $\mathfrak{s}$-space the matrices are now more localized compared to the $\mathfrak{n}$-space in which they have a high variance. The spaces spanned by the columns of the matrices are orthogonal to each other, that is $\mathcal{Q} \perp \mathcal{W}$.

discussed in Section 2.5. In the context of SPD matrices, the JBLD is referred to as the symmetric Stein divergence or the (square of the) log-determinant metric [27]. In the sequel we adopt the notation $\delta_{\mathrm{s}}(X, Y)$.

Motivated by the above, we formulate SSA using a new cost function based on $\delta_{\mathrm{s}}$ (cf. Eqs. (3) and (4)):

$$\mathcal{L}(\widehat{B}) = \sum_{i=1}^{N} \delta_{\mathrm{s}}^2(\widehat{\Sigma}_i^{\mathfrak{s}}, \mathbb{I}) = \sum_{i=1}^{N} \delta_{\mathrm{s}}^2(Q^\top \tilde{\Sigma}_i Q, \mathbb{I}) = \sum_{i=1}^{N} \left[ \log\det\left(\frac{1}{2}(\widehat{\Sigma}_i^{\mathfrak{s}} + \mathbb{I})\right) - \frac{1}{2}\log\det(\widehat{\Sigma}_i^{\mathfrak{s}}) \right], \tag{7}$$

where $Q = \left(\mathbb{I}_D^m \widehat{B}\right)^\top$ and $\tilde{\Sigma}_i = Z\Sigma_i Z^\top$ are the matrices whitened with $Z = \bar{\Sigma}^{-1/2}$, $\bar{\Sigma} = \underset{\Sigma \in \mathcal{S}_D^+}{\operatorname{argmin}} \delta_{\mathrm{s}}^2(\Sigma_i, \Sigma)$, the mean of $\Sigma_i$ w.r.t. $\delta_{\mathrm{s}}$. $\mathcal{S}_D^+$ denotes the set of all $D \times D$ SPD matrices.

## 2.3 A geometric interpretation

By replacing the cost function with one based on the symmetrized divergence we gain not only the beneficial properties of the symmetric divergence, but also new insight into the problem of SSA. First, note that in Eq. (7) the problem is ultimately framed in terms of distances between SPD matrices. This suggests adopting a geometric perspective, whereby the notion of stationarity is captured by the dispersion of the matrices $\Sigma_i$. In this view, the assumption that the covariance matrices of stationary signals do not vary much between epochs translates to them having small distances between them.

An illustration of this idea is presented in Fig. 1. In this figure, the matrices $\Sigma_i$ are seen as points on the SPD matrix manifold $\mathcal{S}_D^+$. The goal of our method is to find transformations $Q$ and $W$ that map the matrices onto two separate manifolds of lower dimension - the stationary and non-stationary space, respectively. The matrices in the stationary space will exhibit small variation,

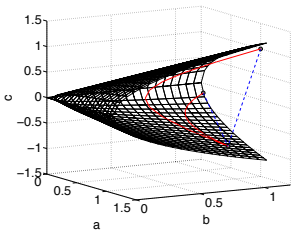


Figure 2: Comparison between Euclidean (blue straight dashed lines) and Riemannian (red curved solid lines) distances measured between points of the space $\mathcal{S}_2^+$.

while the non-stationarities will be captured in the non-stationary space where the variation of the matrices will be greater. The transformations $Q$ and $W$ may be chosen to be orthogonal to each other, producing well separated $\mathfrak{s}$- and $\mathfrak{n}$-spaces. That is, $W \in \mathcal{Q}^\perp$ for $\mathcal{Q} = \operatorname{span} Q$ (likewise $Q \in \mathcal{W}^\perp$), where $\perp$ denotes the orthogonal complement.

Put formally, our objective is to find a rank-$m$ transformation matrix $Q \in \mathbb{R}^{D\times m}$ which maps $\Sigma_i \in \mathcal{S}_D^+$ to $\widehat{\Sigma}_i^{\mathfrak{s}} \in \mathcal{S}_m^+$ for $m < D$ such that the log-determinant distance between the compressed centered matrices $\widehat{\Sigma}_i^{\mathfrak{s}} = Q^\top \tilde{\Sigma}_i Q$ and their mean, which for the centered matrices is $\mathbb{I}$, is minimized. Note that the space spanned by the columns of $Q$ is of importance, and not the specific columns themselves. So, we may optimize $Q$ over the Grassmann manifold [12], $\mathcal{G} = \{\operatorname{span}(Q) : Q \in \mathbb{R}^{D\times m},\ Q^\top Q = \mathbb{I}\}$, the set of all $m$-dimensional linear subspaces of $\mathbb{R}^{D\times D}$.

In practice, for optimization we employ a Riemannian trust-regions method described in Absil et al. [2] and implemented efficiently in Boumal et al. [8].

## 2.4 A generic geometric formulation

Given the strong relation between the log-determinant metric and the AIRM [7], a natural progression is to incorporate the AIRM into the cost function. To understand why it would be beneficial to use the AIRM it is necessary first to briefly discuss the geometry of the SPD matrix manifold.

When equipped with the Frobenius inner product $\langle A, B\rangle_{\mathcal{F}} = \operatorname{tr}(A^\top B)$, the set $\mathcal{S}_n^+$ of SPD matrices of size $n \times n$, belongs to a Euclidean space. In this case, similarity between SPD matrices can be measured simply by using the Euclidean distance derived from the Euclidean norm. This is readily seen in the following example for $2 \times 2$ SPD matrices. A matrix $A \in \mathcal{S}_2^+$ can be written as $A = \left[\begin{smallmatrix} a & c \\ c & b \end{smallmatrix}\right]$ with $ab - c^2 > 0$, $a > 0$ and $b > 0$. Then matrices in $\mathcal{S}_2^+$ can be represented as points in $\mathbb{R}^3$ and the constraints can be plotted as a convex cone whose interior is populated by the SPD matrices (see Fig. 2). In this representation, the Euclidean geometry of symmetric matrices then implies that distances are computed along straight lines.

Despite its simplicity, the Euclidean geometry has several drawbacks and is not always well suited for SPD matrices [13, 3]. For example, due to an

artifact referred to as the *swelling effect* [3], for a task as simple as averaging two matrices, it may occur that the determinant of the average is larger than any of the two matrices. Another drawback, illustrated in Fig. 2 and documented by Fletcher et al. [13], is the fact that this geometry forms a non-complete space. Hence, in this Euclidean space interpolation between SPD matrices is possible, but extrapolation may produce indefinite matrices, leading to uninterpretable solutions.

An efficient alternative which addresses these issues is to consider the space of SPD matrices as a curved space, namely a Riemannian manifold. Of the possible Riemannian distances, the AIRM, due to its favorable mathematical properties, is widely used in many applications (see, for example [13, 24]). It is defined for any $X, Y \in \mathcal{S}_D^+$ as:

$$\delta_{\mathrm{r}}^2(X,Y) = \left\| \log\left(X^{-1/2} Y X^{-1/2}\right) \right\|_{\mathbb{F}}^2, \tag{8}$$

where $\log(\cdot)$ is the matrix logarithm and $\|X\|_{\mathbb{F}}^2 = \mathrm{tr}\left(X^\top X\right)$ is the Frobenius norm.

In the curved space, the geodesics between matrices obtained by the AIRM are computed on curved lines as illustrated in Fig. 2 for the space $\mathcal{S}_2^+$. Symmetric matrices with null and infinite eigenvalues (i.e., those which lie on the boundary of the convex cone, but not in it) are both at an infinite distance from any SPD matrix on the manifold (*within* the cone).

So, let us also consider a cost function of the same form defined w.r.t. the AIRM.[2] A general expression for our geometry-aware SSA (gaSSA) is then:

**Definition 2.1 (gaSSA)** *gaSSA w.r.t. a metric $\delta$ is defined as*

$$\widehat{Q} = \underset{Q \in \mathcal{G}(D,m)}{\operatorname{argmin}} \sum_i \delta^2\left(Q^\top \tilde{\Sigma}_i Q, \mathbb{I}\right), \tag{9}$$

where $\tilde{\Sigma}_i = Z \Sigma_i Z^\top$ are the matrices whitened with $Z = \bar{\Sigma}^{-1/2}$ and $\bar{\Sigma} = \underset{\Sigma \in \mathcal{S}_D^+}{\operatorname{argmin}} \delta^2(\Sigma_i, \Sigma)$ is the matrix mean w.r.t. $\delta$. We note that this cost function is similar in spirit to the one in Harandi et al. [14] and can be considered an unsupervised version of it. In the next section we will show that the need for matrix whitening can be alleviated.

The Euclidean gradient w.r.t. $Q$ of the cost function, used for the optimization, can be found in [15].

## 2.5 Symmetries and invariance properties

We now discuss the symmetries of our optimization problem and the invariance properties of our chosen metrics. These properties will enable us to significantly

[2] Other metrics such as the Euclidean metric or the log-Euclidean metric [3] may also be used.

simplify our problem. For brevity we will state the results in terms of the AIRM, but the same holds true for the log-determinant metric.

Our key observation stems from the fact that $\delta_{\mathrm{r}}$ and $\delta_{\mathrm{s}}$ are invariant to congruent transformations of the form $X \mapsto P^H X P$ for $P \in GL_D(\mathbb{C})$ and $P^H = \overline{P}^\top$ is the conjugate transpose [7]. However, since our discussion is limited to real matrices, we have

$$\delta_{\mathrm{r}}^2 (X, Y) = \delta_{\mathrm{r}}^2 \left( P^\top X P, P^\top Y P \right). \tag{10}$$

for $X, Y \in \mathcal{S}_D^+$ and a real-valued invertible matrix $P$. This is a crucial point since the whitening matrix $Z$ and, more importantly, the *mixing matrix* $A$, act on the covariance matrices in this way.

**Proposition 2.2** *Let $\mathbf{\Lambda} = \{\Lambda_i\}_{i=1}^n$ for $\Lambda_i \in \mathcal{S}_n^+$ be a set of SPD matrices of size $n \times n$ and let $\Sigma_i = A \Lambda_i A^\top$ for some real-valued invertible matrix $A$. Denoting the Riemannian mean of $\mathbf{\Lambda}$ by $\bar{\Lambda}$, the Riemannian mean $\bar{\Sigma}$ of the set $\mathbf{\Sigma} = \{\Sigma_i\}_{i=1}^n$ is given by $\bar{\Sigma} = A \bar{\Lambda} A^\top$.*

*Proof.* The Riemannian mean of the set $\mathbf{\Lambda}$ is defined as $\bar{\Lambda} = \underset{\Lambda \in \mathcal{S}_n^+}{\operatorname{argmin}} \sum_i \delta_{\mathrm{r}}^2 (\Lambda_i, \Lambda)$. Using the congruence invariance (Eq. (10)) we have $\delta_{\mathrm{r}}^2 (\Lambda_i, \Lambda) = \delta_{\mathrm{r}}^2 \left( \Sigma_i, A \Lambda A^\top \right)$ and the result follows. ∎

Using the above we obtain several useful equivalence relations.

**Corollary 2.3** *The following expressions are equivalent:*

$$\delta_{\mathrm{r}}^2 \left( \tilde{\Sigma}_i, \mathbb{I} \right) = \delta_{\mathrm{r}}^2 \left( \Sigma_i, \bar{\Sigma} \right) = \delta_{\mathrm{r}}^2 \left( \Gamma_i, \bar{\Gamma} \right), \tag{11}$$

*where $\Gamma_i$ is the covariance matrix of the* unmixed *sources in the $i$-th epoch.*

We have essentially shown that both the whitening operation and the mixing matrix $A$ do not affect the distance between the covariance matrices of the original unmixed signals. We can then re-write our optimization problem as

$$\begin{aligned} \widehat{Q} &= \underset{Q \in \mathcal{G}(D,m)}{\operatorname{argmin}} \sum_i \delta^2 \left( Q^\top \tilde{\Sigma}_i Q, \mathbb{I} \right) = \underset{Q \in \mathcal{G}(D,m)}{\operatorname{argmin}} \sum_i \delta^2 \left( Q^\top A \Gamma_i A^\top Q, Q^\top A \bar{\Gamma} A^\top Q \right) \\ &= \underset{Q' \in \mathcal{G}(D,m)}{\operatorname{argmin}} \sum_i \delta^2 \left( Q'^\top \Gamma_i Q', Q'^\top \bar{\Gamma} Q' \right). \end{aligned} \tag{12}$$

One may remark that $A^\top Q$ no longer has orthonormal columns and so does not belong to the Grassmann manifold. Indeed this is true. The final transition is due to the observation that the solution to our optimization problem is not unique. Rather, since we are interested in recovering the stationary subspace and not the exact sources themselves, the solution is invariant to any transformation (e.g., subspace scaling and rotation) acting within each of the $\mathfrak{s}$- and $\mathfrak{n}$-spaces separately. Furthermore, we have chosen $W \in \mathcal{Q}^\perp$, and so the $\mathfrak{s}$-space is orthogonal to the $\mathfrak{n}$-space. Now, choosing orthogonal bases within

each of the subspaces we may restrict ourselves to *orthogonal* mixing matrices $A$ and find a transformation $Q$ which lies in the Grassman manifold.

The final result is quite remarkable. First, it shows that our problem is essentially agnostic to the mixing matrix.Secondly, it eliminates the need to pre-whiten the matrices. This is useful in situations where it is not appropriate to whiten the data in advance, for example, when working with nearly ill-conditioned matrices. Furthermore, this eliminates any error that may be caused by inaccurate estimation of the matrix mean.

In conclusion, we have two variations of gaSSA given in the first and last terms of Eq. (12). The difference between the two is whether or not the input covariance matrices are whitened. Our analysis shows that whitening does not improve performance, and may in fact lead to a degradation of the results in certain cases. So, we claim that it is in general preferable not to whiten the matrices. In terms of the chosen metric, we do not expect a significant difference when using $\delta_{\mathrm{r}}$ vs. $\delta_{\mathrm{s}}$. In the following section we will present experimental evidence to support these claims.

# 3 Experimental results

In this section we present experimental results on synthetic and data taken from a real BCI experiment. We compare the performance of gaSSA to the existing SSA and investigate the effects of matrix whitening and the choice of metric.

## 3.1 Toy data

For our first experiment we generated data following the SSA model as a mixture of stationary and non-stationary sources. To generate non-stationarity in the data we used a slightly modified version of the scheme provided in the SSA toolbox [22] and detailed in its user manual. Here we bring only a brief description:

The elements of the mixing matrix $A$ are chosen uniformly from the range $[-0.5, 0.5]$ and its columns are normalized to 1. The distribution of the $\mathfrak{s}$-sources is constant over all epochs, namely $s^{\mathfrak{s}}(t) \sim \mathcal{N}(0, \Lambda^{\mathfrak{s}})$. In the SSA toolbox, $\Lambda^{\mathfrak{s}}$ is taken to be the identity matrix, however, we choose $\Lambda^{\mathfrak{s}}$ to be a random matrix of the form $\Lambda^{\mathfrak{s}} = B\Gamma B^{\top}$ for an orthogonal matrix $B$ and diagonal matrix $\Gamma$.

The $\mathfrak{n}$-sources are correlated with the $\mathfrak{s}$-sources, and for the $i$-th epoch $\tau_i = [t_0(i), \ldots, t_0(i) + T]$ they are given by $s^{\mathfrak{n}}(t) = C_i s^{\mathfrak{s}}(t) + Y^{\mathfrak{n}}(t)$ for $\quad t \in \tau_i$, where $C_i \in \mathbb{R}^{(D-m)\times m}$ and $Y^{\mathfrak{n}}(t) \sim \mathcal{N}(\mu_i, \Lambda_i^{\mathfrak{n}})$. The covariance matrices $\Lambda_i^{\mathfrak{n}}$ are generated for each epoch in the same way as $\Lambda^{\mathfrak{s}}$.

So, the covariance matrix of the (unmixed) sources in the $i$-th epoch may be written as

$$\Lambda_i = \operatorname{cov}\left(\begin{bmatrix} s^{\mathfrak{s}}(t) \\ s^{\mathfrak{n}}(t) \end{bmatrix}\right) = \begin{bmatrix} \Lambda^{\mathfrak{s}} & (C_i\Lambda^{\mathfrak{s}})^{\top} \\ C_i\Lambda^{\mathfrak{s}} & C_i\Lambda^{\mathfrak{s}}C_i^{\top} + \Lambda_i^{\mathfrak{n}} \end{bmatrix}. \tag{13}$$

Using the data generated by the scheme above we compared the performance of our method to that of SSA. We used the AIRM and log-determinant metric both with and without matrix whitening.

As a performance measure we used the distance between the estimated $\mathfrak{n}$-space $\widehat{A}^{\mathfrak{n}}$ and the true $\mathfrak{n}$-space $A^{\mathfrak{n}}$. This is owing to the fact that, as discussed in von Bünau et al. [31], the $\mathfrak{n}$-space and $\mathfrak{s}$-sources are identifiable, while the $\mathfrak{s}$-space and $\mathfrak{n}$-sources are not. To illustrate this, note that to be stationary, the $\mathfrak{s}$-sources must consist strictly of stationary sources, while the $\mathfrak{n}$-sources will remain non-stationary even if they include a mixture of stationary signals. The distance between sub-spaces is computed using $\delta_{\mathcal{G}}$, the metric on the Grassmann manifold [1]. Shortly, this metric is based on the principal angles between the two spaces.

We generated 50 epochs of length $T = 250$ for several values of $D$ and $m$. For each pair $(D, m)$ we conducted the experiment 25 times. At each iteration the optimization procedure was restarted 5 times with different initial guesses and the transformation matrix which obtained the lowest cost was selected.

Due to length restrictions we describe only the results for $D = 19$ and $m = 12$. Other choices of parameters exhibited the same behavior. The $\mathfrak{n}$-space errors were identical for $\delta_{\mathrm{r}}$ and $\delta_{\mathrm{s}}$ $(0.0067 \pm 0.0001)$. They were consistently lower than SSA $(0.0115 \pm 0.0004)$. For both metrics the scheme without matrix whitening performed slightly better than those including whitening $(0.0063 \pm 0.0001$ vs. $0.0067 \pm 0.0001)$. These results are consistent with the analysis of Section 2.5.

### 3.2 Brain-computer interface

Next, we applied our method to data taken from the BCI competition IV dataset II. This dataset contains motor imagery (MI) EEG signals affected by eye movement artifacts. It was collected in a multi-class setting, with the subjects performing more than 2 different MI tasks. However, as in Lotte and Guan [21], we evaluate our algorithms on two-class problems by selecting only signals of left- and right-hand MI trials.

We applied the same pre-processing as described in Lotte and Guan [21]. EEG signals were band-pass filtered in $8 - 30$ Hz, using a $5^{\text{th}}$ order Butterworth filter. For each trial, we extracted features from the time segment located from 0.5s to 2.5s after the cue instructing the subject to perform MI.

The data was initially divided into two parts: a training data set and a test data set. Similarly to von Bünau et al. [31] the first 20% of the test trials were set aside for adaptation. The aim of the adaptation part is to mitigate any non-stationarities between the test and the training session. We then learned the $\mathfrak{s}$-space in an *unsupervised* manner over the training and adaptation part. As before, our method was reinitialized 5 times and the transformation attaining the lowest cost was chosen.

The performance was measured by means of the classification rate on the test set. We used the following naive classifier, referred to as *minimum distance to the mean* (MDM) in Barachant et al. [5]: Using the labels of the training set, we compute the mean (in the $\mathfrak{s}$-space) for each of the two classes. Then, we classify the compressed covariance matrices in the test set according to their distance to the class means; each matrix is assigned the class to which it is closer.

Table 1: Classification accuracy for $m = 10$ (top) and $m = 14$ (bottom) $\mathfrak{s}$-sources. Best results are highlighted in boldface. (w) and (nw) signify that matrix whitening was / was not performed.

| subject # | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | avg |
|---|---|---|---|---|---|---|---|---|---|---|
| gaSSA $\delta_{\mathrm{r}}$ (w) | 48.96 | 55.65 | 63.48 | 58.22 | **60.69** | 56.26 | 68.09 | 71.30 | 48.70 | 59.04 |
| gaSSA $\delta_{\mathrm{r}}$ (nw) | 73.91 | 59.13 | **91.3** | **69.3** | 57.67 | **66.43** | 58.70 | **93.91** | **86.09** | **72.94** |
| gaSSA $\delta_{\mathrm{s}}$ (w) | 49.83 | 55.65 | 63.57 | 61.43 | 58.9 | 55.7 | **68.17** | 72.17 | 47.34 | 59.2 |
| gaSSA $\delta_{\mathrm{s}}$ (nw) | **74.13** | **60** | **91.3** | 68.52 | 57.03 | 63.96 | 59.48 | **93.91** | 85.99 | 72.7 |
| SSA | 47.74 | 57.16 | 60 | 57.61 | 56.02 | 54.17 | 67.3 | 75.89 | 53.82 | 58.86 |
| gaSSA $\delta_{\mathrm{r}}$ (w) | 65.02 | 53.91 | 72.35 | 65.22 | **62.61** | 59.13 | 72.93 | 77.39 | 77.39 | 67.33 |
| gaSSA $\delta_{\mathrm{r}}$ (nw) | **74.4** | **58.26** | 88.74 | 67.83 | 60 | **65.22** | 68.91 | **94.78** | **87.83** | **74** |
| gaSSA $\delta_{\mathrm{s}}$ (w) | 65.02 | 53.91 | 72.35 | 65.22 | **62.61** | 58.87 | 72.83 | 77.39 | 76.52 | 67.19 |
| gaSSA $\delta_{\mathrm{s}}$ (nw) | **74.4** | **58.26** | **88.78** | 67.83 | 60 | 63.83 | 68.15 | 93.91 | **87.83** | 73.67 |
| SSA | 67.05 | 54.72 | 66.7 | **68.7** | 56.87 | 55.57 | **73.48** | 75.65 | 77.39 | 66.24 |

The original data is comprised of 22 signals. Since the true number of stationary signals is unknown, we repeated the experiment for several values in the range $m \in [10, 18]$. Due to lack of space, we bring the results only for the values $m = 10$ and $m = 14$, the former being an example of challenging task and the latter an example of an intermediate one. The results for the nine subjects in the dataset are summarized in Table 1. The full results can be found in the supplementary material.

The results show that our method outperforms SSA for most subjects. As predicted, the methods without pre-whitening generally performed better than those which included whitening of the covariance matrices. In this more complex setting, more accurate estimation of the mixing matrix does not guarantee better classification. In terms of the metric, we see that $\delta_{\mathrm{r}}$ and $\delta_{\mathrm{s}}$ perform roughly the same.

# 4 Conclusion

We presented a covariance-based method for unsupervised stationary subspace analysis. The problem was phrased in terms of the distance between matrices and not, as in SSA, using the divergence between probability distributions. Owing to the symmetries of the problem and the invariance properties of the geometries, we derived useful equivalence relations. Experiments on both synthetic and BCI data supported our theoretical analysis and showed that our method outperforms SSA.

# References

[1] P.-A. Absil, R. Mahony, and R. Sepulchre. Riemannian geometry of Grassmann manifolds with a view on algorithmic computation. *Acta Applicandae Mathematica*,

80(2):199–220, 2004.

[2] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2009.

[3] V. Arsigny, P. Fillard, X. Pennec, and N. Ayache. Geometric means in a novel vector space structure on symmetric positive-definite matrices. *SIAM Journal on Matrix Analysis and Applications*, 29(1):328–347, 2007.

[4] A. Banerjee, S. Merugu, I. S. Dhillon, and J. Ghosh. Clustering with Bregman divergences. *The Journal of Machine Learning Research*, 6:1705–1749, 2005.

[5] A. Barachant, S. Bonnet, M. Congedo, and C. Jutten. Multiclass brain–computer interface classification by riemannian geometry. *IEEE Transactions on Biomedical Engineering*, 59(4):920–928, 2012.

[6] A. Barachant, S. Bonnet, M. Congedo, and C. Jutten. Classification of covariance matrices using a riemannian-based kernel for bci applications. *Neurocomputing*, 112:172–178, 2013.

[7] R. Bhatia. *Positive Definite Matrices*. Princeton University Press, 2009.

[8] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre. Manopt, a Matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research*, 15:1455–1459, 2014. URL `www.manopt.org`.

[9] A. Cherian and S. Sra. Riemannian sparse coding for positive definite matrices. In *European Conference on Computer Vision (ECCV)*, pages 299–314, 2014.

[10] A. Cherian, S. Sra, A. Banerjee, and N. Papanikolopoulos. Efficient similarity search for covariance matrices via the Jensen-Bregman log-det divergence. In *2011 IEEE International Conference on Computer Vision (ICCV)*, pages 2399–2406. IEEE, 2011.

[11] I. Cohen and B. Berdugo. Speech enhancement for non-stationary noise environments. *Signal processing*, 81(11):2403–2418, 2001.

[12] A. Edelman, T. A. Arias, and S. T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM Journal on Matrix Analysis and Applications*, 20 (2):303–353, 1998.

[13] T. P. Fletcher, C. Lu, S. M. Pizer, and S. Joshi. Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Transactions on Medical Imaging*, 23(8):995–1005, 2004.

[14] M. Harandi, M. Salzmann, and R. Hartley. From manifold to manifold: geometry-aware dimensionality reduction for spd matrices. In *European Conference on Computer Vision (ECCV)*, pages 17–32, 2014.

[15] I. Horev, F. Yger, and M. Sugiyama. Intrinsic PCA for SPD matrices. In *Asian Conference on Machine Learning (ACML)*, 2015.

[16] A. Hyvärinen, J. Karhunen, and E. Oja. *Independent Component Analysis*, volume 46. John Wiley & Sons, 2004.

[17] H. Kantz and T. Schreiber. *Nonlinear Time Series Analysis*, volume 7. Cambridge University Press, 2004.

[18] B. Kulis, M. A. Sustik, and I. S. Dhillon. Low-rank kernel learning with Bregman matrix divergences. *The Journal of Machine Learning Research*, 10:341–376, 2009.

[19] M. J. Kusner, N. I. Kolkin, S. Tyree, and K. Q. Weinberger. Stochastic covariance compression. *arXiv preprint arXiv:1412.1740*, 2014.

[20] O. Ledoit and M. Wolf. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2):365–411, 2004.

[21] F. Lotte and C. Guan. Regularizing common spatial patterns to improve bci designs: unified theory and new algorithms. *IEEE Transactions on Biomedical Engineering*, 58(2):355–362, 2011.

[22] J. S. Müller, P. v. Bünau, F. C. Meinecke, F. J. Király, and K.-R. Müller. The stationary subspace analysis toolbox. *The Journal of Machine Learning Research*, 12:3065–3069, 2011.

[23] F. Nielsen and R. Nock. Sided and symmetrized Bregman centroids. *IEEE Transactions on Information Theory*, 55(6):2882–2904, 2009.

[24] X. Pennec, P. Fillard, and N. Ayache. A Riemannian framework for tensor computing. *International Journal of Computer Vision*, 66(1):41–66, 2006.

[25] W. Samek, K.-R. Müller, M. Kawanabe, and C. Vidaurre. Brain-computer interfacing in discriminative and stationary subspaces. In *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 2873–2876. IEEE, 2012.

[26] H. Shimodaira. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of Statistical Planning and Inference*, 90(2): 227–244, 2000.

[27] S. Sra. Positive definite matrices and the s-divergence. *arXiv preprint arXiv:1110.1773*, 2011.

[28] M. Sugiyama and M. Kawanabe. *Machine Learning in Non-Stationary Environments*. The MIT Press, 2012.

[29] O. Tuzel, F. Porikli, and P. Meer. Region covariance: A fast descriptor for detection and classification. In *European Conference on Computer Vision*, pages 589–600, 2006.

[30] P. von Bünau, F. C. Meinecke, F. C. Király, and K.-R. Müller. Finding stationary subspaces in multivariate time series. *Physical Review Letters*, 103(21):214101, 2009.

[31] P. von Bünau, F. C. Meinecke, and K.-R. Müller. Stationary subspace analysis. In *Independent Component Analysis and Signal Separation*, pages 1–8. Springer, 2009.

[32] W. Wojcikiewicz, C. Vidaurre, and M. Kawanabe. Stationary common spatial patterns: towards robust classification of non-stationary eeg signals. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 577–580. IEEE, 2011.